

Maximizing Profit in Green Cellular Networks through Collaborative Games

Please, cite this paper as:

Cosimo Anglano, Marco Guazzone, Matteo Sereno,
*“Maximizing Profit in Green Cellular Networks,
through Collaborative Games,”*

Computer Networks, Volume 75, Part A, pp. 260–275, 2014.

DOI: 10.1016/j.comnet.2014.10.003

Publisher: <http://www.sciencedirect.com/science/article/pii/S1389128614003582>

Maximizing Profit in Green Cellular Networks through Collaborative Games^{☆,☆☆}

Cosimo Anglano^{a,*}, Marco Guazzone^a, Matteo Sereno^b

^a*Department of Science and Technological Innovation, University of Piemonte Orientale, Italy*

^b*Department of Computer Science, University of Torino, Italy*

Abstract

In this paper, we deal with the problem of maximizing the profit of *Network Operators* (NOs) of green cellular networks in situations where *Quality-of-Service* (QoS) guarantees must be ensured to users, and *Base Stations* (BSs) can be shared among different operators.

We show that if NOs cooperate among them, by mutually sharing their users and BSs, then each one of them can improve its net profit.

By using a game-theoretic framework, we study the problem of forming *stable* coalitions among NOs. Furthermore, we propose a mathematical optimization model to allocate users to a set of BSs, in order to reduce costs and, at the same time, to meet user QoS for NOs inside the same coalition. Based on this, we propose an algorithm, based on cooperative game theory, that enables each operator to decide with whom to cooperate in order to maximize its profit.

This algorithm adopts a distributed approach in which each NO autonomously makes its own decisions, and where the best solution arises without the need to synchronize them or to resort to a trusted third party.

The effectiveness of the proposed algorithm is demonstrated through a thorough experimental evaluation considering real-world traffic traces, and a set of realistic scenarios. The results we obtain indicate that our algorithm allows a population of NOs to significantly improve their profits thanks to the combination of energy reduction and satisfaction of QoS requirements.

Keywords: Wireless Networks, Profit Maximization, Green Networking, Cooperative Game Theory, Coalition Formation

[☆]Please, cite as: Cosimo Anglano, Marco Guazzone, Matteo Sereno, “Maximizing Profit in Green Cellular Networks through Collaborative Games,” *Computer Networks*, Volume 75, Part A, Pages 260–275, 2014, DOI:10.1016/j.comnet.2014.10.003.

^{☆☆}Link to publisher: <http://www.sciencedirect.com/science/article/pii/S1389128614003582>

*Corresponding author

Email addresses: cosimo.anglano@di.unipmn.it (Cosimo Anglano), marco.guazzone@di.unipmn.it (Marco Guazzone), matteo.sereno@unito.it (Matteo Sereno)

1. Introduction

The increasing consumption of electrical energy is one of the most important issues characterizing modern society because of its effects on climate changes and on the depletion of non-renewable sources. In this scenario, the ICT sector plays a key role, being responsible for about 10% of the world carbon footprint and electrical energy consumption [1, 2].

Reportedly [3, 4], within the ICT sector, the mobile telecommunication industry (and, in particular, cellular networks) is one of the major contributors to energy consumption. This has stimulated the interest towards a new research area called *green cellular networks* [5], that aims at reducing the energy consumption of these communication infrastructures.

From the perspective of a *cellular Network Operator* (NO), the reduction of electrical energy consumption is not only a matter of being “green” and responsible, but also an economically important opportunity. As a matter of fact, it has been argued that nearly half of the total operating expenses of a NO is due to energy costs [3, 4]. Furthermore, a significant part of these costs are due to *Base Stations* (BSs) [6]: indeed, even in the case of little or no activity, a BS can consume more than 90% of its peak energy [4, 7]. Thus, by reducing energy consumption, a NO may sensibly increase its profit.

Consequently, a lot of research effort has been concentrated lately on the reduction of the energy consumed by BSs. Techniques like the design of more energy-efficient hardware equipments, or the use of new energy saving techniques (e.g., *sleep modes* [8] and *cell zooming* [9]) to switch off under-utilized BSs during low traffic periods and to transfer the corresponding load to neighboring cells, have been proposed as possible solutions.

Such techniques, however, must be applied with care so as to maintain *Quality-of-Service* (QoS) guarantees agreed by a NO with its customers, whose violations imply monetary losses for that NO. Specifically, since fewer transmission resources are available at a cell when such energy-efficient techniques are used, bottlenecks may form for those users connected to that cell, who may thus experience QoS levels lower than guaranteed, and in some cases may be even unable to receive service at all. Finally, the use of techniques like cell zooming may cause other problems, such as inter-cell interference and coverage holes [9].

In this paper, we argue that, if NOs cooperate among them by mutually sharing their users and BSs, then each one of them can improve its net profit by either (a) reducing energy costs by switching off its BSs and offloading its users to switched on BSs of other NOs, or (b) increasing its earnings by attracting users from other NOs, or by relying on BSs of other NOs to accept more users than what could do by working alone.

Obviously, it is unreasonable to expect that each NO is willing to unconditionally cooperate with the other ones regardless the benefits it receives. Such a cooperation arises indeed only if suitable benefits result from it, and if the risks of monetary losses are kept within acceptable limits.

In this paper, we devise a decision algorithm that provides a set of NOs with suitable means to decide whether to cooperate with other NOs, and if so with whom to cooperate. Our algorithm is based on game-theoretic techniques, where the process of establishing cooperation among the NOs is modeled as a *cooperative game with transferable utility* [10] (in particular, as a *hedonic game* [11], whereby each NO bases its decision on its own

preferences).

More specifically, we propose a game-theoretic framework to study the problem of forming *stable* coalitions among NOs, and a mathematical optimization model to allocate users to a set of BSs, in order to reduce costs and, at the same time, to meet user QoS for NOs inside the same coalition. We achieve our goal by devising a *hedonic shift algorithm* to form stable coalitions that allows each NO to autonomously and selfishly decide whether to leave the current coalition to join a different one or not on the basis of the net profit it receives for doing so.

In our approach, each NO pays for the energy consumed to serve each user, whether it belongs to it or to another NO, but receives a payoff (computed as discussed later) for doing so. We prove that the proposed algorithm converges to a *Nash-stable* set of disjoint coalitions [12], whereby no NO can benefit to leave the current coalition to join a different one.

Our solution adopts an asynchronous approach in which each NO autonomously makes its own decisions, and where the best solution arises without the need to synchronize them or to resort to a trusted third party. As a consequence, the solution we propose can be readily implemented in a distributed fashion.

To demonstrate the effectiveness of the algorithm we propose, we carry out a thorough experimental evaluation considering real-world traffic traces, and a set of realistic scenarios. The results we obtain indicate that our algorithm allows indeed a population of NOs to significantly improve their profits thanks to the combination of energy reduction and satisfaction of QoS requirements.

The contributions of this paper can be summarized as follows:

- we consider the problem of maximizing operators' profit in green cellular networks;
- we model the problem as a cooperative game with transferable utility;
- we devise a distributed algorithm enabling operators to find the coalition maximizing their profits under stability concerns;
- we show its effectiveness through experimental analysis in realistic scenarios;
- we assess the impact of energy price and user population on the profits attained by operators.

The rest of this paper is organized as follows. In Section 2, we describe the system under study and we present the problem addressed in this paper. In Section 3, we present the cooperative game-theoretic framework we use to study the problem of coalition formation and the hedonic shift algorithm we design to form stable coalitions. In Section 4, we show results from an experimental evaluation to show the effectiveness of the proposed approach. In Section 5, we provide an overview of related works. Finally, in Section 6, we conclude the paper and present an outlook on possible future extensions.

2. System Model and Problem Definition

2.1. System Model

We consider an area served by a set $\mathcal{N} = \{1, \dots, N\}$ of NOs, whose BSs fully cover that area and whose coverage overlaps (as typically happens in urban areas [9, 13, 14]). To keep the notation simple, we assume that in the area of interest there is only one BS per NO, so in the rest of this paper, we will use the terms BS and NO interchangeably (the extension of the model to support multiple BSs per NO is straightforward).

Each BS i is characterized by its maximum downlink transmission capacity C_i , and by its power consumption $W_i(n_i)$ that, as argued in [7, 15, 16], is linearly dependent on the number of users it is serving, that is:

$$W_i(n_i) = \alpha_i + \beta_i n_i \quad (1)$$

where α_i (the *static term*) is the load-independent power consumption (which is usually known from the specifications of the BS, and typically accounts for about 90% of the total consumption [4, 7]), and $\beta_i n_i$ (the *dynamic term*) is the load-dependent power consumption (that can be determined by linear regression from real power measurements [16]).

Each NO i provides network connectivity to a set \mathcal{U}_i of *customers* (hereafter also referred to as *users*). Each user $j \in \mathcal{U}_i$ is characterized by its required QoS, quantified by the minimum downlink data rate D_j it requests, and the actual downlink data rate d_j it gets from the network.

Each user $j \in \mathcal{U}$ (where $\mathcal{U} = \bigcup_{i=1}^N \mathcal{U}_i$) can connect to any BS in the system regardless of the NO who owns the BS (i.e., it can connect to a BS that belongs to the NO to which it is subscribed or it can roam on the BS of another one). This can be accomplished by using techniques like *cell wilting* and *blossoming* [17]. However, the aggregate allocated data rate to users connected to BS i cannot exceed its capacity C_i , that is:

$$\sum_{j \in \mathcal{U}_i} d_j \leq C_i. \quad (2)$$

We assume that the number of users receiving service from a BS i varies over time, and is described by the *load profile* curve $\ell_i(t)$ of that BS that expresses, as function of time, the percentage of the maximum number of users M_i that can receive service by BS i when each user j is allocated its entire desired data rate D_j . It then follows that, if all the users have the same data rate requirement (i.e., $D_j = D$ for all $j \in \mathcal{U}_i$), then $M_i = C_i/D$. Conversely, if users are heterogeneous, then M_j is estimated as $M_j = C_i/\bar{D}$, where \bar{D} is the weighted average of the data rates requested by users, i.e., $\bar{D} = \sum_{j \in \mathcal{U}_i} p_j D_j$, where p_j is the probability that user j arrives at BS i .

An example of a typical daily load profile is depicted in Figure 1, where the x -axis represents the time (in hours) and the y -axis is the normalized load of the BS [18]. For instance, if at a given time t , $M_i = 10$ and $\ell_i(t) = 0.8$, the number $n_i(t)$ of user of BS i at time t is $n_i(t) = 0.8 \cdot 10 = 8$.




Figure 1: A typical daily load profile $\ell_i(\cdot)$ of a single BS i .

2.2. Problem Definition

Given the system characterized as above and a particular area of interest, each NO seeks to maximize its *net profit* (i.e., the difference between its revenues and costs) in the presence of a time-varying population of users in this area.

The *net profit rate* P_i of NO i (i.e., the profit it makes per unit of time) can be expressed as

$$P_i = \sum_{j \in \mathcal{U}_i} R_{i,j} - \left[W_i(n_i) E_i + \sum_{j \in \mathcal{U}_i} L_{i,j}(d_j) \right] \quad (3)$$

where $R_{i,j}$ is the revenue rate generated by user j on BS i , E_i is the electricity cost rate of BS i , and $L_{i,j}(d_j)$ is the penalty rate incurred by NO i if user j receives a downlink rate d_j lower than its QoS value D_j , which is given by the following loss function:

$$L_{i,j}(d_j) = \left(1 - \frac{d_j}{D_j} \right) R_{i,j} \quad (4)$$

Thus, $L_{i,j}(d_j)$ is zero if the QoS of the user is completely satisfied (i.e., $d_j = D_j$), and linearly increases until $R_{i,j}$ as the assigned data rate d_j decreases (so that an NO gets no revenue from those customers that receive no service).

If the NOs in the area of interest cooperate among them (i.e., they share their users and BSs) then each NO i can maximize the corresponding value of P_i by acting on the various terms of Eq. (3) as follows:

- it can attempt to reduce $W_i(n_i)$ by offloading (some of) its users to the BSs of other operators so that its BS can be switched off entirely (by exploiting *sleep modes* [8]) or only in part (by relying on *cell zooming* [9]);
- it can attempt to increase $R_{i,j}$ either by attracting users from other NOs, so that it can better amortize its energy cost $W_i(n_i)$, or by relying on BSs of other NOs to accept users that, if working alone, it could not serve without violating Eq. (2), thus incurring into a (possibly high) penalty rate $L_{i,j}$.

It is evident that, to exploit these opportunities, each NO must be willing to cooperate with (at least some of) the other ones. In the following section, we will characterize the conditions under which such a cooperation is not only possible, but also sought by these NOs.

3. The Coalition Formation Game

As discussed before, cooperation is the key to increase profit. However, it is unreasonable to assume that a NO is willing to unconditionally cooperate with the other ones regardless of the benefits it receives. As a matter of fact, the acceptance of users roaming from other NOs is beneficial only if the additional revenue they bring outweighs the costs and the possible penalties they induce. Furthermore, the offloading of users to other NOs makes sense only if a suitable revenue results from this operation for the off-loader.

To cooperate, a set of NOs must first form a *coalition*, i.e., they all must agree to share their own BSs and users among them. Given a set of NOs, however, there can be many different coalitions that can be formed, each one differing from the other ones in terms of the structure (i.e., the identity of each member) and/or of the profit it brings to their members.

In order to join a coalition, a NO must indeed find it *profitable*, i.e., it must be sure that the profit it earns by joining the coalition is no worse of the one it obtains by working alone. Furthermore, in order to be sure that this profit is not ephemeral, a NO must seek other properties that guarantee the suitability of a coalition, namely:

- *Stability*: a coalition is *stable* if none of its participants finds that it is more profitable to leave it (e.g., to stay alone or to join another coalition) rather than cooperating with the other ones. Lack of stability causes possible monetary losses for the following reasons:
 - a NO that has joined a coalition with the expectation of receiving users roaming from other NOs is penalized if, after switching on a BS on which to accommodate these users, these NOs leave the coalition;
 - a NO that has accepted more users than those it can serve without incurring into a penalty, expecting to use the BSs of other NOs to accommodate them, is penalized if these NOs leave the coalition.
- *Fairness*: when joining a coalition, a NO expects that the resulting profits are fairly divided among participants. As an unfair division leads to instability, a fair profit allocation strategy is mandatory.

From these considerations, it clearly follows that a way must be provided to each NO to decide whether to participate to a coalition or not and, if so, which one among all the possible coalitions is worth joining.

In this paper, we address this issue by modeling the problem of coalition formation as a *coalition formation cooperative game with transferable utility* [10, 19], where each NO cooperates with the other ones in order to maximize its net profit rate, and by devising an

algorithm to solve it. By using our algorithm, the various NOs can make their decisions concerning coalition membership.

In the rest of this section, we first set the coalition formation problem in the game-theoretic framework (Section 3.1), then we present an algorithm to form stable coalitions among NOs (Section 3.2), and finally we present an optimization model to allocate users to a set of BSs, in order to reduce costs and, at the same time, to meet user QoS for NOs inside the same coalition (Section 3.3).

3.1. Characterization

Our coalition formation algorithm is based on a *hedonic game* [11], a class of *coalition formation cooperative games* [10, 19] where each NO acts as a selfish agent and where its preferences over coalitions depend only on the composition of that coalition. That is, NOs prefer being in one coalition rather than in another one solely based on who else is in the coalitions they belong.

Formally, given the set $\mathcal{N} = \{1, 2, \dots, N\}$ of NOs (henceforth also referred to as the *players*), a *coalition* $\mathcal{S} \subseteq \mathcal{N}$ represents an agreement among the NOs in \mathcal{S} to act as a single entity.

At any given time, the set of players is partitioned into a *coalition partition* Π , that we define as the set $\Pi = \{\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_l\}$, where $\mathcal{S}_k \subseteq \mathcal{N}$ ($k = 1, \dots, l$) is a disjoint coalition such that $\bigcup_{k=1}^l \mathcal{S}_k = \mathcal{N}$ and $\mathcal{S}_j \cap \mathcal{S}_k = \emptyset$ for $j \neq k$. Given a coalition partition Π , for any NO $i \in \mathcal{N}$, we denote as $\mathcal{S}_\Pi(i)$ the coalition to which i is participating.

Each coalition \mathcal{S} is associated with its *coalition value* $v(\mathcal{S})$, that we define as the *net profit rate* of that coalition, that is:

$$v(\mathcal{S}) = R(\mathcal{U}_\mathcal{S}) - Q(\mathcal{U}_\mathcal{S}) - K(\mathcal{S}) \quad (5)$$

where:

- $R(\mathcal{U}_\mathcal{S})$ is the *coalition revenue rate*, corresponding to the sum of revenue rates of individual users $j \in \mathcal{U}_\mathcal{S}$ (where $\mathcal{U}_\mathcal{S} = \bigcup_{i \in \mathcal{S}} \mathcal{U}_i$ is the joint user population of the NOs belonging to \mathcal{S});
- $Q(\mathcal{U}_\mathcal{S})$ is the *coalition load cost rate*, and is computed by minimizing the costs resulting from serving the users in $\mathcal{U}_\mathcal{S}$ using all the resources provided by the NOs belonging to \mathcal{S} (we discuss this in Section 3.3);
- $K(\mathcal{S})$ is the *coalition formation cost rate*, that takes into account the cost incurred by players to establish and maintain the coalition (e.g., the costs for system reconfiguration to enable user migration and handover across NOs). In this paper, we assume $K(\mathcal{S})$ to be proportional to the coalition size, and we define it as:

$$K(\mathcal{S}) = \begin{cases} \sum_{i \in \mathcal{S}} K_i, & |\mathcal{S}| > 1, \\ 0, & \text{otherwise.} \end{cases} \quad (6)$$

where K_i is the coalition formation cost rate for NO i .

Obviously, each NO $i \in \mathcal{S}$ must receive a fraction $x_i(\mathcal{S})$ of the coalition value, that we call the *payoff* of i in \mathcal{S} . Our game is conceived in such a way to form coalitions in which NOs get payoffs as high as possible, without violating the fairness requirement, so that stability is achieved. Thus, a *payoff allocation rule* must be specified in order to compute the payoffs of each coalition member in such a way to ensure fairness in the division of payoffs.

To this end, we use the *Shapley value* [20], a payoff allocation rule that is based on the concept of *marginal contribution* of players (i.e., the change in the worth of a coalition when a player joins to that coalition), such that the larger is the contribution provided by a player to a coalition, the higher is the payoff allocated to it.¹ This means that, in a given coalition, some “more-contributing” NOs will be rewarded by other “less-contributing” NOs to encourage them to join the coalition. More specifically, the Shapley value $\phi_i(v)$ of player i is defined as:

$$\phi_i(v) = \sum_{\mathcal{S} \subseteq \mathcal{N} \setminus \{i\}} \frac{|\mathcal{S}|!(N - |\mathcal{S}| - 1)!}{N!} (v(\mathcal{S} \cup \{i\}) - v(\mathcal{S})) \quad (7)$$

where the sum is over all subsets \mathcal{S} not containing i (the symbol “ \setminus ” denotes the set difference operator), and the symbol “ $!$ ” denotes the factorial function.

It is worth noting that we rely on the Shapley value for its interesting properties. Nevertheless, other payoff allocation rules can be used and our work is general enough to support them.

To set up the coalition formation process, we need to define, for each NO i , a *preference relation* \succeq_i that NO i can use to order and compare all the possible coalitions it may join. Formally, this corresponds to define a complete, reflexive, and transitive binary relation over the set of all coalitions that NO i can form (see [12]).

Specifically, for any NO $i \in \mathcal{N}$ and given $\mathcal{S}_1, \mathcal{S}_2 \subseteq \mathcal{N}$, the notation $\mathcal{S}_1 \succeq_i \mathcal{S}_2$ means that NO i prefers being a member of \mathcal{S}_1 over \mathcal{S}_2 or at least i prefers both coalitions equally. In our coalition formation game, for any NO $i \in \mathcal{N}$, we use the following preference relation:

$$\mathcal{S}_1 \succeq_i \mathcal{S}_2 \Leftrightarrow u_i(\mathcal{S}_1) \geq u_i(\mathcal{S}_2), \quad (8)$$

where $\mathcal{S}_1, \mathcal{S}_2 \subseteq \mathcal{N}$ are any two coalitions that contain NO i (i.e., $i \in \mathcal{S}_1$ and $i \in \mathcal{S}_2$), and u_i is a preference function defined for any NO i as follows:

$$u_i(\mathcal{S}) = \begin{cases} x_i(\mathcal{S}), & \mathcal{S} \notin h(i), \\ -\infty, & \text{otherwise.} \end{cases} \quad (9)$$

where $x_i(\mathcal{S})$ is the payoff received by NO i in \mathcal{S} by means of Eq. (7), and $h(i)$ is a *history* set where NO i stores the identity of the coalitions that have been already evaluated so that we avoid generating twice the same candidate coalition (a similar idea for pruning already considered coalitions has also been used in previously published work, such as in [22]).

¹More specifically, we use the *Aumann-Dr ze* value [21], which is an extension of the Shapley value for games with coalition structures.

Thus, according to Eq. (9), each NO prefers to join the coalition that provides the larger payoff.

The strict counterpart of \succeq_i , denoted by \succ_i , is defined by replacing \geq with $>$ in Eq. (8), and implies that i strictly prefers being a member of \mathcal{S}_1 over \mathcal{S}_2 .

3.2. The Algorithm for Coalition Formation

In this section, we present an algorithm for coalition formation that allows the NOs to take distributed decisions for selecting which coalitions to join at any point in time. This algorithm is based on the following *hedonic shift rule* (see [22]):

Definition 1. *Given a coalition partition $\Pi = \{\mathcal{S}_1, \dots, \mathcal{S}_h\}$ on the set \mathcal{N} and a preference relation \succ_i , any NO $i \in \mathcal{N}$ decides to leave its current coalition $\mathcal{S}_{\Pi}(i) = \mathcal{S}_l$, for $1 \leq l \leq h$, to join another one $\mathcal{S}_k \in \Pi \cup \emptyset$, with $\mathcal{S}_k \neq \mathcal{S}_l$, if and only if $\mathcal{S}_k \cup \{i\} \succ_i \mathcal{S}_l$, that is if its payoff in the new coalition exceeds the one it is getting in its current coalition. Hence, $\{\mathcal{S}_l, \mathcal{S}_k\} \rightarrow \{\mathcal{S}_l \setminus \{i\}, \mathcal{S}_k \cup \{i\}\}$.*

This shift rule (that we denote as “ \rightarrow ”) provides a mechanism through which any NO can leave its current coalition $\mathcal{S}_{\Pi}(i)$ and join another coalition \mathcal{S}_k , given that the new coalition $\mathcal{S}_k \cup \{i\}$ is strictly preferred over $\mathcal{S}_{\Pi}(i)$ through any preference relation that the NOs are using. This rule can be seen as a selfish decision made by a NO to move from its current coalition to a new one, *regardless of the effects of this move on the other NOs*.

Using the hedonic shift rule, we design a distributed hedonic coalition formation algorithm for NOs as presented in Algorithm 1.

The basic idea of the algorithm is to have each NO i search, asynchronously with respect to the other NOs, the state space of possible coalitions it may join, and for each one of them, evaluate whether it is preferable (according to the corresponding \succ_i relation) to remain in its current coalition, or to join it. Whenever a NO decides to move from a coalition to another one, it updates its history set $h(i)$ by appending the coalition it is leaving, so that the same coalition is not visited twice during the coalition space search. A NO iterates the actions listed in Algorithm 1 until no more hedonic shift rules are possible. It is worth noting that the asynchronicity of our algorithm makes it suitable to be executed, for instance, when new users arrive to NOs, thus making it able to adapt to environmental changes.

Let us explain in detail how Algorithm 1 works. The algorithm takes as parameters the global state *state*, storing the current shared coalition partition Π_c , and the identity i of the calling NO (initially there are no coalitions, i.e., $\Pi_c = \Pi_0 = \{\{1\}, \{2\}, \dots, \{N\}\}$).

At each execution of the algorithm, NO i initializes its history set h and other auxiliary variables (lines 2–3), and then enters a loop that is executed until no more hedonic shift rules can be performed from the last coalition partition considered by i .

In each loop iteration, NO i retrieves the current coalition partition, and generates all the possible hedonic shifts until no more of them are possible. Given the distributed nature of the algorithm, we postulate the use of suitable distributed space management algorithms (e.g. [23, 24]).

Then, after acquiring a lock to gain exclusive access to the shared state (line 5) in order to ensure its atomic update (by means of a suitable distributed mutual exclusion algorithm

Algorithm 1 The Coalition Formation Algorithm for NOs

```

1: procedure COALITIONFORMATION( $state, i$ )
2:    $h \leftarrow \emptyset$ 
3:    $\mathcal{S}_{\text{best}} \leftarrow \emptyset$ 
4:   repeat
5:     LOCK( $state$ )
6:      $\Pi_c \leftarrow \text{GETCURRENTPARTITION}(state)$ 
7:      $\mathcal{S}_{\text{cur}} \leftarrow \mathcal{S}_{\Pi_c}(i)$ 
8:      $\mathcal{S}_{\text{best}} \leftarrow \mathcal{S}_{\text{cur}}$ 
9:     for all  $\mathcal{S} \in (\Pi_c \setminus \{\mathcal{S}_{\text{cur}}\}) \cup \emptyset$  and  $\mathcal{S} \notin h$  do
10:       $\mathcal{S}_{\text{new}} \leftarrow \mathcal{S} \cup \{i\}$ 
11:       $x_{\mathcal{S}_{\text{best}}} \leftarrow \text{COMPUTEPAYOFF}(\mathcal{S}_{\text{best}}, i)$  ▷ See Eq. (5) and Eq. (7)
12:       $x_{\mathcal{S}_{\text{new}}} \leftarrow \text{COMPUTEPAYOFF}(\mathcal{S}_{\text{new}}, i)$  ▷ See Eq. (5) and Eq. (7)
13:      if  $x_{\mathcal{S}_{\text{new}}} > x_{\mathcal{S}_{\text{best}}}$  then ▷ See Eq. (8) and Eq. (9)
14:         $\mathcal{S}_{\text{best}} \leftarrow \mathcal{S}_{\text{new}}$ 
15:      end if
16:    end for
17:    if  $\mathcal{S}_{\text{best}} \neq \mathcal{S}_{\text{cur}}$  then
18:       $\mathcal{S} \leftarrow \mathcal{S}_{\text{cur}} \setminus \{i\}$ 
19:       $\mathcal{T} \leftarrow \mathcal{S}_{\text{best}} \setminus \{i\}$ 
20:      UPDATEHISTORY( $h, \mathcal{S}$ )
21:       $\Pi_{\text{best}} \leftarrow (\Pi_c \setminus \{\mathcal{S}_{\text{cur}}, \mathcal{T}\}) \cup \{\mathcal{S}, \mathcal{S}_{\text{best}}\}$ 
22:      SETCURRENTPARTITION( $state, \Pi_{\text{best}}$ )
23:    end if
24:    UNLOCK( $state$ )
25:  until  $\mathcal{S}_{\text{best}} = \mathcal{S}_{\text{cur}}$ 
26: end procedure

```

[23]), NO i iteratively evaluates all the possible coalitions it can form from its current coalition partition, to look for the one with the higher payoff.

To do so, given its current coalition partition Π_c , for each coalition $\mathcal{S}_k \in \Pi_c \cup \emptyset$ (not present in its history set and different from its current one $\mathcal{S}_{\Pi_c}(i)$), NO i applies the hedonic shift rule and evaluates its preference against the current coalition $\mathcal{S}_{\Pi_c}(i)$ (lines 9–16).

If a coalition \mathcal{S}_k with the higher payoff is found (lines 17–23), NO i adds to its history set h the coalition $\mathcal{S}_{\Pi_c}(i) \setminus \{i\}$ it is leaving, and updates the partition set by updating both \mathcal{S}_k (that now contains also i) and $\mathcal{S}_{\Pi_c}(i)$ (that now does not contain i anymore).

Then, after releasing the exclusive lock to the shared state (line 24), NO i repeats the above steps (lines 5–24) to look for a better coalition, in case some other NO has meanwhile modified the shared state by changing the coalition partition.

Eventually, if no other better coalition is found, NO i terminates the execution of the algorithm (line 25), until a new instance is run again.

The convergence of the proposed algorithm during the hedonic coalition formation phase is guaranteed as follows:

Proposition 1 (Convergence). *Starting from any initial coalition structure Π_0 , the proposed algorithm always converges to a final partition Π_f .*

Proof. The coalition formation phase can be mapped to a sequence of shift operations. That is, according to the hedonic shift rule, every shift operation transforms the current partition Π_c into another partition Π_{c+1} . Thus, starting from the initial step, the algorithm yields the following transformations:

$$\Pi_0 \rightarrow \Pi_1 \rightarrow \dots \rightarrow \Pi_c \rightarrow \Pi_{c+1} \quad (10)$$

where the symbol “ \rightarrow ” denotes the application of a shift operation. Every application of the shift rule leads to a coalition partition that has not been previously visited (i.e., a new coalition partition). Thus, the number of transformations performed by the shift rule is finite (at most, it is equal to the number of partitions, that is the Bell number) and hence the sequence in Eq. (10) will always terminate and converge to a final partition Π_f . \square

The stability of the final partition Π_f resulting from the convergence of the proposed algorithm can be addressed by using the following stability definition (see [12] for details).

Definition 2. A coalition partition $\Pi = \{\mathcal{S}_1, \dots, \mathcal{S}_l\}$ is *Nash-stable* if $\forall i \in \mathcal{N}$, $\mathcal{S}_{\Pi}(i) \succeq_i \mathcal{S}_k \cup \{i\}$ for all $\mathcal{S}_k \in \Pi \cup \emptyset$.

It is worth noting that Nash-stability captures the notion of stability with respect to movements of single NOs (i.e., no NO has an incentive to unilaterally deviate).

For the hedonic coalition formation phase of the proposed algorithm, we can prove the following result:

Proposition 2 (Nash-stability). *Any final partition Π_f resulting from Algorithm 1 is Nash-stable.*

Proof. We prove it by contradiction. Assume that the final partition Π_f is not Nash-stable. Consequently, there exists a NO $i \in \mathcal{N}$ and a coalition $\mathcal{S}_k \in \Pi_f \cup \emptyset$ such that $\mathcal{S}_k \cup \{i\} \succ_i \mathcal{S}_{\Pi_f}(i)$. Then, NO i will perform a hedonic shift operation and hence $\Pi_f \rightarrow \Pi'_f$, where Π'_f is the new coalition partition resulting after the hedonic shift operation. This contradicts the assumption that Π_f is the final outcome of our algorithm. \square

It is worth to point out that Nash-stability also implies the so called *individual-stability* [12]. A partition $\Pi = \{\mathcal{S}_1, \dots, \mathcal{S}_l\}$ is *individually-stable* if it does not exist a NO $i \in \mathcal{N}$ and a coalition $\mathcal{S}_k \in \Pi \cup \emptyset$ such that $\mathcal{S}_k \cup \{i\} \succ_i \mathcal{S}_{\Pi}(i)$ and $\mathcal{S}_k \cup \{i\} \succeq_j \mathcal{S}_k$ for all $j \in \mathcal{S}_k$, i.e., if no NO can benefit by moving from its coalition to another existing (possibly empty) coalition while not making the members of that coalition worse off. Thus, we can conclude that our algorithm always converges to a partition Π_f which is both Nash-stable and individually stable.

Given the NP-completeness of the problem of finding a Nash-stable partition [25], the computational cost of our algorithm can become quite large when the number of NOs increases (indeed, in the worst case, it is bounded by the N^{th} Bell number, where N is the number of NOs). In these cases, we can reduce the computational cost by having each NO check for the Nash-stability of its current coalition partition (this check takes polynomial time [25]), and execute the algorithm only if Nash-stability no longer holds true.

3.3. Computation of the Optimal Coalition Load Cost

The algorithm presented in the previous section requires the computation of the value $v(\mathcal{S})$ of any coalition \mathcal{S} that each NO i may possibly join, that in turn requires the computations of coalition load cost rate $Q(\mathcal{U}_{\mathcal{S}})$ (see Eq. (5)). To compute $Q(\mathcal{U}_{\mathcal{S}})$, we need in turn to determine, for the coalition \mathcal{S} , the optimal data rate allocation (i.e., the allocation of users that minimizes the costs of NOs).

To this end, we define a *Mixed Integer Linear Program* (MILP) modeling the problem of allocating a set $\mathcal{U}_{\mathcal{S}}$ of users onto a set \mathcal{S} of BSs so that the overall cost rates of NOs in \mathcal{S} are minimized. The resulting optimization model is shown in Figure 2, where we use the same notation introduced in Section 2 (however, to ease readability, we denote with \mathcal{U} the user set, i.e., we drop the dependence from \mathcal{S}).

In the optimization model we use the following decision variables:

- $u_{i,j}$, which is a binary variable that is equal to 1 if user j is allocated to BS i ;
- $d_{i,j}$, which is a real variable representing the downlink data rate allocated to user j by BS i ;
- b_i , which is a binary variable that is equal to 1 if BS i is switched on.

The objective function $Q(\mathcal{U})$ (see Eq. (11a)) represents the cost rates incurred by the coalition of NOs for serving users in \mathcal{U} , and is defined as the sum of the costs due to the power absorbed by the BSs that are switched-on, and of those due to QoS violations (if any).

The resulting optimal user allocation is bound to the following constraints:

$$\text{minimize } Q(\mathcal{U}) = \sum_{i \in \mathcal{S}} \left[b_i W_i \left(\sum_{j \in \mathcal{U}} u_{i,j} \right) E_i + \sum_{j \in \mathcal{U}} L_{i,j}(d_{i,j}) \right] \quad (11a)$$

subject to

$$\sum_{j \in \mathcal{U}} d_{i,j} \leq C_i, \quad i \in \mathcal{S}, \quad (11b)$$

$$\sum_{i \in \mathcal{S}} u_{i,j} = 1, \quad j \in \mathcal{U}, \quad (11c)$$

$$\sum_{j \in \mathcal{U}} u_{i,j} \leq b_i U, \quad i \in \mathcal{S}, \quad (11d)$$

$$d_{i,j} \leq u_{i,j} D_j, \quad i \in \mathcal{S}, j \in \mathcal{U}, \quad (11e)$$

$$d_{i,j} \in \mathbb{R}^*, \quad i \in \mathcal{S}, j \in \mathcal{U}, \quad (11f)$$

$$u_{i,j} \in \{0, 1\}, \quad i \in \mathcal{S}, j \in \mathcal{U}, \quad (11g)$$

$$b_i \in \{0, 1\}, \quad i \in \mathcal{S}. \quad (11h)$$

Figure 2: The user-to-BS allocation optimization model.

- Eq. (11b) ensures that the capacity of a switched-on BS is not exceeded;
- Eq. (11c) imposes that each user is served by exactly one BS;
- Eq. (11d) states that only BSs that are switched on can serve users; the purpose of this constraint is to avoid that a user is served by a BS that will be switched off;
- Eq. (11e) imposes that each user obtains at most the requested data rate by the serving BS;
- Eq. (11f), Eq. (11g), and Eq. (11h) define the domain of decision variables $d_{i,j}$, $u_{i,j}$, and b_i , respectively.

As can be noted from the definition of $Q(\mathcal{U})$, the solution of the optimization problem at a specif instant of time requires the knowledge of the number of users present in each BS i at that time. In general, however, such number is not constant, but it varies over time according to the corresponding load profile $\ell_i(t)$.

In order to compute the number of users of BS i at time t from ℓ_i , we proceed as follows: first, as typically done in the literature [26, 27], we discretize $\ell_i(t)$ by splitting the time axis into uniform disjoint sub-intervals $[\tau, \tau + \Delta t)$ of length Δt time units (where Δt is the *discretization step*). Then, we approximate the (normalized) load of each subinterval as a constant value set to the peak load of that subinterval. For instance, the result of the discretization of the load profile of Figure 1 with Δt of 1 hour is depicted in Figure 3. In this figure, the time-horizon of one day (i.e., 24 hours) is split into several subintervals




Figure 3: Discretization of the load profile of Figure 1 with a time-horizon of 1 day and $\Delta t = 1$ hour (vertical segments represent subintervals bounds and horizontal segments are peaks inside subintervals).

$[\tau, \tau + \Delta t)$ of length $\Delta t = 1$ hour, where $\tau = 0, 1, \dots, 23$. Each subinterval is delimited by vertical dotted segments, while every horizontal solid red segment is the peak inside each subinterval, that we will use as an approximation of the (normalized) load inside the subinterval.

4. Experimental Evaluation

In order to assess the ability of our algorithm of increasing the net profits for a population of NOs, we perform a set of experiments in which we consider a variety of realistic scenarios and real-world traffic data. In these experiments we vary, in a controlled way, various input parameters of the algorithm, namely the cost of energy, the QoS requirements of users, and the discretization step of the traffic profile curve, so that we are able to assess the impact of each one of them on the performance of the algorithm. The results we collect, discussed in this section, demonstrate the ability of our algorithm of yielding significant increases of the net profit achieved by a set of NOs in all the scenarios we consider.

To perform such experiments, we develop an ad-hoc simulator written in C++ and interfaced with CPLEX [28] to solve the various instances of the optimization model presented in Section 3.3.

4.1. Experimental Setup

We consider a system configuration comprising five NOs, each one owning a single BS. Without loss of generality, we assume that all the BSs are identical in terms of capacity and energy consumption. More specifically, we set $C_i = 100$ Mbps, $\alpha_i = 0.551$ kW, and $\beta_i = 0.00146$ kW for $i = \{1, \dots, 5\}$ (these last two values have been taken from [16]). We also assume that all NOs incur in the same coalition cost rate, i.e., $K_i = 0.01$ \$/hour $\forall i \in \{1, \dots, 5\}$.

Furthermore, we assume that each BS has its own load profile, that differs from those of the other ones. The load profiles we consider in our experiments, reported in Figs. 4a–4e, have been obtained from real-world data [29] consisting of normalized cellular traffic

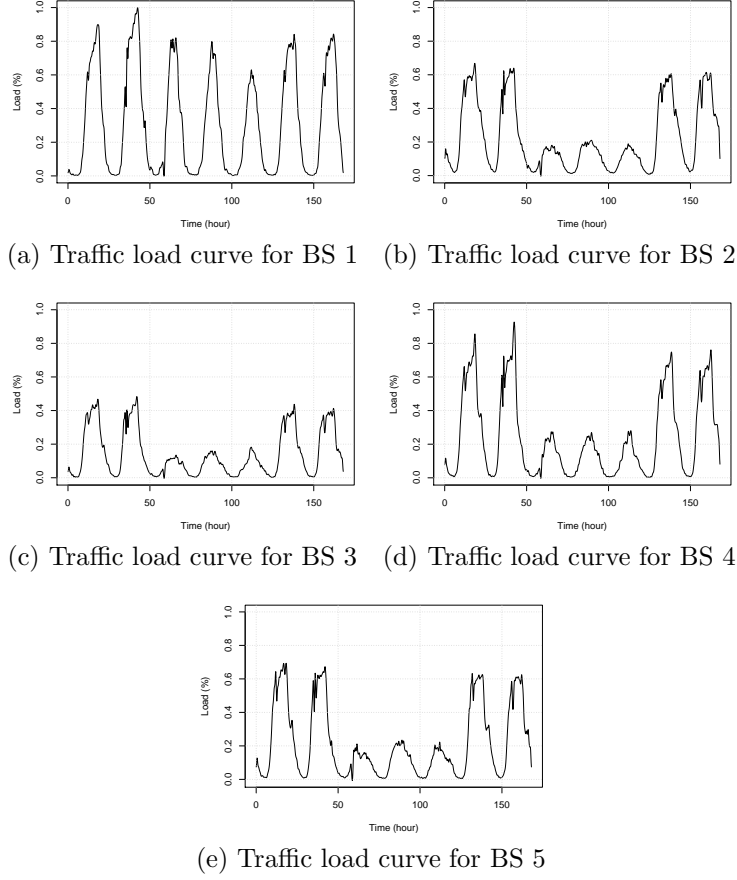


Figure 4: Traffic load curves for the BSs of the experimental scenarios.

collected, with a resolution of 30 minutes, in a metropolitan urban area during one week, and been already used in similar studies [13, 30]. Specifically, each load profile curve $\ell_i(\cdot)$ of Figure 4 has been obtained by fitting a periodic cubic spline to the traffic data related to BS i . We characterize each load profile by computing the overall load during the entire week as $\Gamma_i = \int_0^{168} \ell_i(t) dt, \forall i \in \{1, \dots, 5\}$, as well as the average hourly load as $\bar{\ell}_i = \Gamma_i/168$, that are reported in Table 1 (the upper integration limit corresponds to the number of hours corresponding in a week).

4.2. Experimental Results

To evaluate the performance of our algorithm, we compute the relative net profit increment RP_i attained by each NO i , that is defined as:

$$RP_i = \sum_{k=1}^A \frac{x_i(\mathcal{S}^{(k)})}{P_i^{(k)}} - 1$$

where $A = \lceil 168/\Delta t \rceil$ is the number of executions of the algorithm, while $x_i(\mathcal{S}^{(k)})$ and $P_i^{(k)}$ correspond to the payoff received by NO i at the k -th algorithm execution and the profit it

Table 1: Load characteristics of each BS $i \in \mathcal{N}$.

BS	Γ_i	$\bar{\ell}_i$
1	53.07	0.316
2	37.17	0.221
3	23.95	0.143
4	40.26	0.240
5	36.59	0.218

would attain if it worked alone (computed as in Eq. (3)), respectively. Note that, in general, several Nash-stable coalitions may result at each algorithm execution; in these cases, $x_i(\mathcal{S}^{(k)})$ is computed as the average of the payoffs yielded by all the Nash-stable coalitions that may form.

To explain the results, we also compute, for each NO i , two additional quantities, namely:

- ON_i , the ratio of the number of times that BS i is switched on after an execution of the algorithm over the total number of algorithm executions;
- XL_i , the difference between the ratio of the total number of users served by BS i when working in a coalition over the total number of users it would serve if it was working alone, and 1; this quantity corresponds to the relative deviation of load experienced by BS i with respect to the case it works alone, where a positive (negative) value represents an increment (decrement) of load with respect to the case of working alone.

In the rest of this section, we discuss the impact of the electricity price first (Section 4.2.1), then we consider the effects of the heterogeneity of user requirements (Section 4.2.2), then we evaluate the impact of the discretization step width (Section 4.2.3), and finally we conclude with a discussion of our findings (Section 4.2.4).

4.2.1. Impact of Energy Costs

The energy cost E_i obviously impacts on the net profit achieved by an NO i . Intuitively, if energy is expensive, NO i may find it profitable to offload its users to other NOs, so that it can switch off its BS. Conversely, if energy is inexpensive, it may try to attract users from other NOs.

To quantify the effects of energy price on the net profits achieved by NOs, we carry out experiments on a set of scenarios obtained by setting E_i to either $E_{lo} = 0.12$ \$/kWh (which is a typical value of electricity cost in the US [31]) or $E_{hi} = 2 \cdot E_{lo} = 0.24$ \$/kWh. Because of space constraints, we discuss only the results corresponding to four of the 32 (i.e., 2^5) distinct scenarios resulting from the assignment of each cost value to each one of the NOs, as indicated in Table 2.

These four scenarios have been selected since they can be considered representative of two opposite situations that may occur in practice: the first two scenarios (1 and 2) correspond indeed to standard situations where all BSs, being located in the same urban area, pay the

Table 2: Electricity cost in the selected experimental scenarios.

Scenario	Electricity Cost (\$/kWh)				
	NO 1	NO 2	NO 3	NO 4	NO 5
1	0.12	0.12	0.12	0.12	0.12
2	0.24	0.24	0.24	0.24	0.24
3	0.12	0.24	0.24	0.12	0.12
4	0.12	0.12	0.12	0.24	0.24

Table 3: Impact of energy price: net profit increments of the various NOs.

Scenario	Net Profit Increment (%)				
	NO 1	NO 2	NO 3	NO 4	NO 5
1	12.10	22.58	40.78	18.15	20.67
2	39.78	101.56	481.44	72.33	90.03
3	15.31	106.52	504.60	24.41	28.37
4	15.18	30.11	57.08	77.14	96.24

same energy price, while the last two scenarios (3 and 4) correspond to possible near-future scenarios where different BSs can draw energy produced by either fossil fuel or renewable sources [5] and, as such, pay different prices. In these scenarios, all users require the same minimum downlink bandwidth (i.e., $D_j = 10$ Mbps) and generate the same revenue rate $R_j = 0.07$ \$/hour (this value is based on the 1 GB/month *Share-Everything* plan from *Verizon Wireless* [32]). Furthermore, we assume that each NO executes an instance of Algorithm 1 at every hour (i.e., $\Delta t = 1$ hour).

Table 3 shows the net profit increment RP_i for each NO i , while Table 4 shows the corresponding ON_i and XL_i values.

Let us start with scenarios 1 and 2. As can be seen from the corresponding rows in Table 3, NO 1 and NO 3 achieve the lowest and the highest net profit increase, respectively, in both scenarios. The corresponding ON_i and XL_i values provide an explanation of these facts. Indeed, while NO 3 is able to switch off its BS most of the times (ON_3 is no larger than 26%), thus reducing its energy cost, NO 1 keeps it switched on for 90% of the times or more, thus incurring into an high energy cost. This, in turn, is due to the characteristics of the respective load profiles: while NO 1 most of the times has too many clients to be able to offload all of them to other NOs, NO 3 is in the opposite situation. As a consequence, NO 1 is able to accept users coming from other NOs (XL_1 is larger than 60%), but the extra revenues they generate is in large part elided by its energy costs. Finally, the profit

Table 4: Impact of energy price: ON and XL values (in %) for the various NO.

Scenario	NO 1		NO 2		NO 3		NO 4		NO 5	
	ON_1	XL_1	ON_2	XL_2	ON_3	XL_3	ON_4	XL_4	ON_5	XL_5
1	92.48	60.44	46.85	33.45	19.79	-30.96	22.03	-47.55	22.70	-42.04
2	90.60	60.43	42.84	26.00	25.89	-9.80	14.02	-63.85	23.38	-32.16
3	89.88	62.45	8.63	-84.90	9.52	-75.98	41.67	13.41	50.00	35.47
4	94.05	76.82	39.29	23.78	48.21	56.80	13.99	-75.27	5.65	-87.18

Table 5: Minimum downlink data rates and revenue rates for the different user classes.

User Class	D_j (Mbps)	Traffic Type	R_j (\$/hour)
Base	0.0122	Voice	0.0175
Standard	0.384	3GPP	0.035
Premium	10	HSPA	0.07

Table 6: Impact of user heterogeneity: net profit increments of the various NOs.

Scenario	Net Profit Increment (%)				
	NO 1	NO 2	NO 3	NO 4	NO 5
1	6.23	10.96	19.84	9.46	10.77
2	16.90	33.94	78.29	28.02	33.82
3	7.95	36.06	81.43	12.48	14.84
4	7.88	14.65	28.36	30.11	36.05

increments of the other NOs (namely, 2, 4, and 5) fall in between the above two extremes: they are indeed able to offload their users more often than NO 1 (they have a lower load), but not as often as NO 3 (their load is higher), as indicated by Table 4.

This phenomenon occurs also in scenarios 3 and 4, where we observe that the largest net profit increases are achieved by the NOs that are associated with the highest energy costs (NOs 2 and 3 in scenario 3, and NO 4 and 5 in scenario 4). As indicated by Table 4, these NOs are indeed often able to offload all their users to other BSs (such a frequency depends on the respective load profiles) so that they can frequently switch off their BSs, and the cost savings they achieve are significant given their higher energy costs. Again, as in scenarios 1 and 2, NO 1 gets the lower profit increase, since its ability to switch its BS off remains lower than the other players.

4.2.2. Impact of User Heterogeneity

In the previous set of experiments, all the users were assumed to request the same minimum downlink data rate. However, in realistic settings, it is reasonable to expect that users with different requirements and revenues co-exist in the same area.

In order to study the impact of the composition of the user population on the ability of our algorithm to yield satisfactory results, we carry out a set of experiments in which, for each one of the scenarios listed in Table 2, users are partitioned in equal proportions into three classes, each one characterized by a different values of the minimum downlink data rate and revenue rate (as reported in Table 5). As in Section 4.2.1, we assume that each NO executes an instance of Algorithm 1 at every hour (i.e., $\Delta t = 1$ hour).

Table 6 shows the net profit increment RP_i for each NO i . As for the case of homogeneous users, we observe that in scenarios 1 and 2, the lowest and the highest net profit increases are achieved by NO 1 (i.e., the NO with the heaviest loaded BS) and NO 3 (i.e., the NO with the lightest loaded BS), respectively, while, in scenarios 3 and 4, the largest net profit increments are reached by the NOs that are associated with the highest energy costs (i.e., NOs 2 and 3 in scenario 3, and NOs 4 and 5 in scenario 4). The explanation of this fact is the same given for the homogeneous users case, so we do not repeat it here.

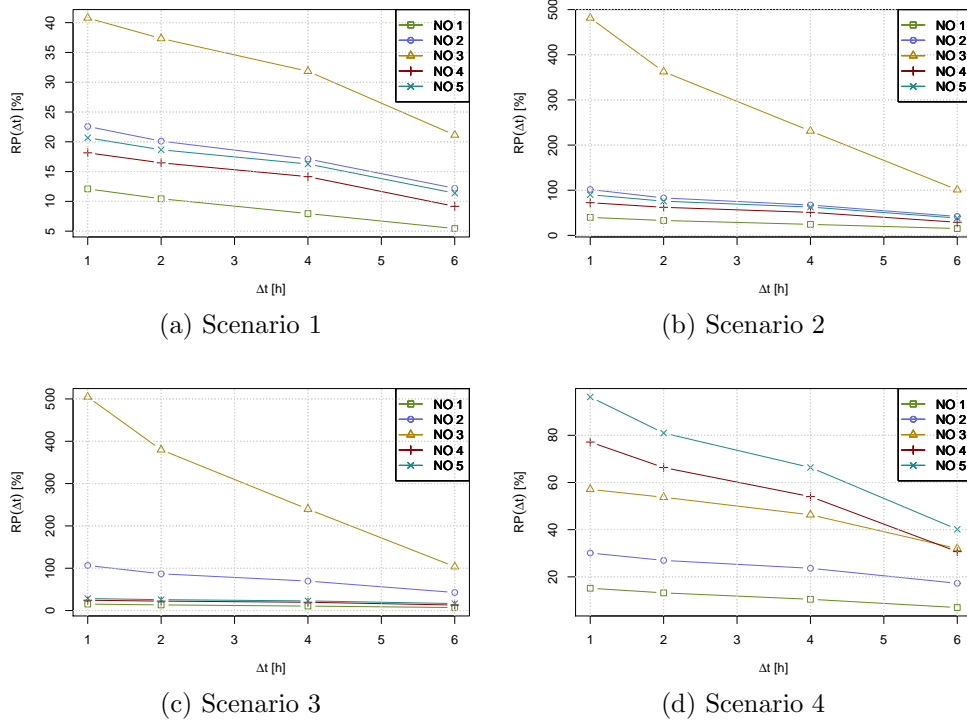


Figure 5: Impact of Δt : RP values for the various NOs and for $\Delta t = 1, 2, 4, 6$.

Furthermore, while the composition of user population has no appreciable effects on the choices taken by individual NOs, it has an evident effect on the net profit increase attained by each NO. As can be indeed seen from Table 6, the RP values are lower than those achieved by each NO in the same scenario when users are homogeneous (reported in Table 3). This is not unexpected, given the lower average revenue rate brought by each user (0.0356 \$/hour versus 0.07 \$/hour). However, we also note that these increases remain significant.

4.2.3. Impact of Δt

The final set of results we comment are concerned with the behavior of the algorithm when the discretization step Δt increases, i.e., when its frequency of activation decreases. This parameter has an evident impact on the solutions computed by the algorithm, since the larger its value, the coarser the approximation of the load value used to compute the optimal coalition load cost $Q(\cdot)$, and the lower the impact of the coalition formation cost $K(\cdot)$.

To quantify this impact, we run several experiments in which we progressively increase Δt . Because of space constraints, here we discuss only the results obtained for $\Delta t = 2, 4, 6$ hours for the same scenarios and user population considered in Section 4.2.1, that are reported in Figure 5. The figure shows for each value of Δt (in the x -axis) the corresponding RP values of the various NOs (denoted as $RP(\Delta t)$, in the y -axis). Also, for the sake of comparability, in the same figure, we report the RP values obtained for $\Delta t = 1$ hour.

As can be seen from Figure 5, the larger Δt , the lower the net profit increase attained

by each NO for a given scenario with respect to when $\Delta t = 1$.

As previously pointed out for the case of $\Delta t = 1$, the energy cost plays an important role in the achieved net profit increments. This still holds for larger value of Δt since, as can be noted in Figure 5, the decrease of the RP values, for increasing values of Δt , is as much larger as the higher is the energy cost. For instance, this can be observed by comparing the RP values obtained in scenarios 1 (see Figure 5a) and 2 (see Figure 5b), whereby the electricity cost double. Here, we can note that as Δt increases, the RP values drops by a factor of about 2 in scenario 1, and by a factor of nearly 2.5 but that can reach a factor of about 4.8 for NO 3, in scenario 2. A similar situation happens in scenarios 3 and 4, where the NOs that experience the higher decrease of their RP values are the ones with associated the higher electricity cost, namely NOs 2 and 3 in scenario 3 (see Figure 5c) and NOs 4 and 5 in scenario 4 (see Figure 5d).

4.2.4. Discussion

From the results obtained in the experimental evaluation we can conclude that:

- energy costs greatly influence the coalition formation and the net profit increments achieved by NOs;
- NOs with higher energy costs are more motivated to join a coalition since they can offload their users to NOs with lower energy costs, thus allowing to switch off their BSs and hence to achieve a higher net profit increment;
- NOs with heavier load are motivated to join a coalition as well, since they can host users of other NOs, thus amortizing their energy costs;
- the composition of user population has no effect on the choices made by the various NOs, while – if the revenue associated with each user is directly proportional to its minimum downlink data rate – it has a strong impact on the net profit increment;
- the discretization step Δt has a significant impact on the coalition formation process, as the larger its value, the lower the net profit increment;
- when choosing the discretization step Δt , one have to also take into account the impact of the energy cost as, the larger is Δt the higher is the (negative) impact of the energy cost on the achieved net profit increment.

5. Related Works

The problem of increasing the profit of NOs in cellular wireless networks has been already studied in the literature, where several papers on this topic have been published.

However, to the best of our knowledge, the problem of forming multiple stable coalitions of BSs, in order to reduce energy consumption and to increase the profit of different and selfish NOs, has never been tackled before. In our work, we pursue this problem by proposing

a novel approach based on mathematical optimization and on the coalition formation game theory.

Much of current research focuses indeed on energy saving techniques as a way to reduce NO costs [5]. For instance, works like [33, 34, 35, 36] use optimization techniques and traffic profile patterns to determine when and where to switch off BSs, while those like [9] use cell-zooming techniques to adaptively adjust the cell size according to traffic load and to possibly switch off inactive cells. Other works, like [37], focus on techniques to improve spectrum efficiency in order to enhance the utilization of data subcarriers and thus to better amortize the license costs of frequency bands. These techniques, however, do not consider the cooperation among different NOs, and therefore are unable, unlike our approach, to exploit the advantages brought by such a cooperation. Furthermore, they do not jointly tackle the problems of ensuring QoS to users and of reducing energy consumption.

Approaches attempting to jointly achieve QoS and energy savings have been recently proposed [27, 38, 39, 40].

In [27], a static joint planning and management optimization approach to limit energy consumption (by switching BSs on and off according to the traffic load) while guaranteeing QoS and minimizing NO costs is proposed. This approach, however, is inherently static (it operates at network design time) so, unlike ours, is unable to operate in a dynamic environment.

In [38], the authors present a cooperative game-theoretic approach, in which individual access networks with insufficient resources join to form the grand-coalition in order to satisfy service demands. This proposal is unable to support non trivial scenarios featuring multiple NOs, a time-varying number of connected users, the energy consumption and the costs due to coalition formation and operation, that may prevent the formation of the grand coalition in favor of smaller and more stable coalitions. Conversely, these scenarios are properly dealt with by our work.

In [39], a game-theoretic approach for the energy-efficient operation of heterogeneous LTE cellular networks, belonging to a single NO, is proposed. This approach, however, does not guarantee the stability of the coalitions that are formed, while stability is a core property of our solution. Furthermore, it is unable to deal with complex scenarios featuring multiple NOs possibly exhibiting different energy prices, and a time-varying population of users. In contrast, our algorithm is able to deal with the above scenarios, and always yield stable coalitions.

In [40], the authors propose a hierarchical dynamic game framework to increase the capacity of two-tier cellular networks by offloading traffic from macro cells to small cells. This work does not take into account the opportunity to selectively switch off underutilized BSs thus offloading the related users to the remaining switched-on BSs. Therefore, our approach can be considered complementary to this one, as it is able to provide a profitable way to select what macro cells to consider before applying the proposed hierarchical game.

6. Conclusion and Future Works

This paper presents a novel dynamic cooperation scheme among a group of cellular NOs to achieve profit maximization. We propose a cooperative game-theoretic framework to study the problem of forming stable coalitions among NOs, and a mathematical optimization model to allocate users to a set of BSs, in order to reduce NO costs and, at the same time, to meet user QoS.

Our solution adopts a distributed approach in which the best solution arises without the need to synchronize the various NOs or to resort to a trusted third party, and such that no NO can benefit by moving from its coalition to another (possibly empty) one.

In the proposed scheme, we model the cooperation among the NOs as a coalition game with transferable utility and we devise a hedonic shift algorithm to form stable coalitions. With our algorithm, each NO autonomously and selfishly decides whether to leave the current coalition to join a different one according to his preference, meanwhile improving its perceived net profit. We prove that the proposed algorithm converges to a Nash-stable partition which determines the resulting coalition structure. Our algorithm can be readily implemented in a distributed fashion, given that each NO can act independently and asynchronously from any other NO in the system. Furthermore, the asynchronicity of our algorithm makes it able to adapt to environmental changes (like new user arrivals).

To prove the effectiveness of our approach, we perform a thorough numerical evaluation by means of trace-based simulation, using realistic scenarios and real-world traffic data. To evaluate the performance of our algorithms, we vary, in a controlled way, the values of its input parameters, like the energy cost and the user heterogeneity.

The future developments of this research is following several directions. Firstly, we want to extend our work to include cellular networks sparsely deployed in wider geographical areas. To do so, we will have to take into account several issues, like the network coverage problem, whereby a BS can be switched off only if a group of neighboring BSs can cover the area it serves.

As a second research direction, we want to explore different variants of our algorithm, especially suited for large cellular networks. Specifically, when the number of BSs increases, the time to convergence of our algorithm may be too long for practical uses, especially for small values of the discretization step. In these cases, it would be better to renounce to the quality of the obtained solution in favor of a more readily available solution. Our current algorithm always provides the best solution (i.e., a Nash-stable partition), but at the cost of visiting, in the worst case, all possible coalitions. To this end, we want to design an *anytime* version of our algorithm (i.e., an algorithm which can return a – possibly suboptimal – solution any time) and we want to compare its performance with the current one.

Finally, we would like to extend our work to include the cooperation between BSs and their users, in order to improve the energy reduction of BSs and the quality of experience of connected users.

References

- [1] Global Action Plan, An inefficient truth, Online: <http://www.globalactionplan.org.uk>, 2007.

- [2] The Climate Group, GeSI, SMART 2020: Enabling the low carbon economy in the information age, Online: http://www.theclimategroup.org/_assets/files/Smart2020Report.pdf, 2008.
- [3] W. Vereecken, W. V. Heddeghem, M. Deruyck, B. Puype, B. Lannoo, W. Joseph, D. Colle, L. Martens, P. Demeester, Power consumption in telecommunication networks: overview and reduction strategies, *IEEE Comm Mag* 49 (2011) 62–69.
- [4] E. Oh, B. Krishnamachari, X. Liu, Z. Niu, Toward dynamic energy-efficient operation of cellular network infrastructure, *IEEE Comm Mag* 49 (2011) 56–61.
- [5] A. D. Domenico, E. C. Strinati, A. Capone, Enabling green cellular networks: A survey and outlook, *Comput Comm* 37 (2014) 5–24.
- [6] G. Auer, V. Giannini, C. Desset, I. Godor, P. Skillermark, M. Olsson, M. Imran, D. Sabella, M. Gonzalez, O. Blume, A. Fehske, How much energy is needed to run a wireless network?, *IEEE Wireless Comm Mag* 18 (2011) 40–49.
- [7] C. Peng, S.-B. Lee, S. Lu, H. Luo, H. Li, Traffic-driven power saving in operational 3g cellular networks, in: *Proc. of the 17th Annual Int. Conference on Mobile Computing and Networking (MobiCom)*, ACM, 2011, pp. 121–132.
- [8] G. Micallef, P. Mogensen, H.-O. Sheck, Cell size breathing and possibilities to introduce cell sleep mode, in: *Proc. of the European Wireless Conference (EW)*, 2010, pp. 111–115.
- [9] Z. Niu, Y. Wu, J. Gong, Z. Yang, Cell zooming for cost-efficient green cellular networks, *IEEE Comm Mag* (2010).
- [10] B. Peleg, P. Sudhölter, *Introduction to the Theory of Cooperative Games*, 2nd ed., Springer Berlin Heidelberg, 2007.
- [11] J. Drèze, J. Greenberg, Hedonic coalitions: Optimality and stability, *Econometrica* 48 (1980) 987–1003.
- [12] A. Bogomolnaia, M. Jackson, The Stability of Hedonic Coalition Structures, *Game Econ Behav* 38 (2002) 201–230.
- [13] M. Ajmone Marsan, L. Chiaraviglio, D. Ciullo, M. Meo, Multiple daily base station switch-offs in cellular networks, in: *Proc. of the 4th Int. Conference on Communications and Electronics (ICCE)*, 2012, pp. 245–250.
- [14] E. Pollakis, R. L. Cavalcante, S. Stańczak, Base station selection for energy efficient network operation with the majorization-minimization algorithm, in: *Proc. of the 13th IEEE Int. Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, 2012.
- [15] M. Deruyck, E. Tanghe, W. Joseph, L. Martens, Characterization and optimization of the power consumption in wireless access networks by taking daily traffic variations into account, *EURASIP J Wireless Comm Networking* 2012 (2012) 1–12.
- [16] J. Lorincz, T. Garma, G. Petrovic, Measurements and Modelling of Base Stations Power Consumption under Real Traffic Loads, *Sensors* 12 (2012) 4281–4310.
- [17] A. Conte, A. Feki, L. Chiaraviglio, D. Ciullo, M. Meo, M. Ajmone Marsan, Cell wilting and blossoming for energy efficiency, *IEEE Wireless Comm Mag* 18 (2011) 50–57.
- [18] D. Willkomm, S. Machiraju, J. Bolot, A. Wolisz, Primary users in cellular networks: A large-scale measurement study, in: *Proc. of the 3rd IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, 2008, pp. 1–11.
- [19] D. Ray, *A Game-Theoretic Perspective on Coalition Formation*, The Lipsey Lectures, Oxford University Press, 2007.
- [20] L. S. Shapley, A Value for n -person Games, in: H. Kuhn, A. Tucker (Eds.), *Contributions to the Theory of Games*, Princeton University Press, 1953, pp. 307–317.
- [21] R. Aumann, J. Drèze, Cooperative games with coalition structures, *Int J Game Theor* 3 (1974) 217–237.
- [22] W. Saad, Z. Han, T. Başar, M. Debbah, A. Hjørungnes, Hedonic coalition formation for distributed task allocation among wireless agents, *IEEE Trans Mobile Comput* 10 (2011) 1327–1344.
- [23] A. Kshemkalyani, M. Singhal, *Distributed Computing: Principles, Algorithms, and Systems*, Cambridge University Press, 2008.
- [24] G. Weiss (Ed.), *Multiagent Systems*, 2nd ed., MIT Press, 2013.
- [25] C. Ballester, NP-completeness in hedonic games, *Game Econ Behav* 49 (2004) 1–30.

- [26] J. Lorincz, A. Capone, D. Begusic, Impact of service rates and base station switching granularity on energy consumption of cellular networks, *EURASIP J Wireless Comm Networking* 342 (2012) 1–24.
- [27] S. Boiardi, A. Capone, B. Sansó, Radio planning of energy-aware cellular networks, *Comput Network* 57 (2013) 2564–2577.
- [28] IBM ILOG CPLEX Optimizer, Online: <http://www.ibm.com/software/integration/optimization/cplex-optimizer>, 2014.
- [29] K. Son, E. Oh, B. Krishnamachari, Normalized cellular traffic during one week, Online: <http://anrg.usc.edu/downloads>, 2011. Dataset No. 18.
- [30] K. Son, H. Kim, Y. Yi, B. Krishnamachari, Base station operation and user association mechanisms for energy-delay tradeoffs in green cellular networks, *IEEE J Sel Area Comm* 29 (2011) 1525–1536.
- [31] U.S. Energy Information Administration, Average retail price of electricity to ultimate customers by end-use sector, Online: <http://www.eia.gov/electricity>, 2013.
- [32] Verizon Wireless, Share Everything plan, Online: <http://www.verizonwireless.com/wcms/consumer/shop/share-everything>, 2014.
- [33] M. Ajmone Marsan, L. Chiaraviglio, D. Ciullo, M. Meo, Optimal energy savings in cellular access networks, in: *Proc. of the IEEE Int. Conference on Communications Workshops (ICC)*, 2009, pp. 1–5.
- [34] H. Ghazzai, E. Yaacoub, M. Alouini, A. Abu-Dayya, Optimized green operation of LTE networks in the presence of multiple electricity providers, in: *Proc. of the IEEE Globecom Workshops (GC Wkshps)*, 2012, pp. 664–669.
- [35] F. Han, Z. S. Z., K. Ray Liu, Energy-efficient base-station cooperative operation with guaranteed QoS, *IEEE Trans Comm* 61 (2013) 3505–3517.
- [36] C. Hasan, E. Altman, J.-M. Gorce, The coalitional switch-off game of service providers, in: *Proc. of the 9th IEEE Int. Conference on Wireless and Mobile Computing, Networking and Communications (WiMob)*, 2013.
- [37] C.-C. Hsu, J. M. Chang, Z.-T. Chou, Z. Abichar, Optimizing spectrum-energy efficiency in downlink cellular networks, *IEEE Trans Mobile Comput* (2013). In press.
- [38] J. Antoniou, I. Koutoussidis, E. Jaho, A. Pitsillides, I. Stavrakakis, Access network synthesis game in next generation networks, *Comput Network* 53 (2009) 2716–2726.
- [39] E. Yaacoub, A. Imran, Z. Dawy, A. Abu-Dayya, A game theoretic framework for energy efficient deployment and operation of heterogeneous LTE networks, in: *Proc. of the 18th IEEE Int. Workshop on Computer Aided Modeling and Design of Communication Links and Networks (CAMAD)*, 2013, pp. 33–37.
- [40] K. Zhu, E. Hossain, D. Niyato, Pricing, spectrum sharing, and service selection in two-tier small cell networks: A hierarchical dynamic game approach, *IEEE Trans Mobile Comput* (2013). In press.